

propagating changes to the database objects based on the top flavor and the at least one other flavor.

13. (New) The method of claim 12, wherein the database objects include one of tables, rows, columns, data containers, records, fields, key columns, and key fields.

REMARKS

By this amendment, claims 1-13 are pending, in which claims 1-3, 5, and 9 are amended, and claims 11-13 are added. No new matter is introduced (see, e.g., FIGs. 2-5, and p. 7:16-20, and discussion thereof in Applicants' Specification).

The present Office Action rejected claims 1-10 under 35 U.S.C. § 103(a) based on *Goldring* (US 5,440,735) in view of *Susuki et al.* (US 6,003,067).

The rejection of claims 1-8 is respectfully overcome, because the applied references fail to teach or suggest the features recited in the claims. For example, independent claim 1, as amended, recites:

wherein the first copy of the table and the second copy of the table have at least one **non-overlapping relational database column**.

By employing a non-overlapping relational database column, advantageously, "communication cost in replication because less data is transferred between the master site 600 and the client site 620" is reduced (Specification, p. 17:18-20). As admitted by the present Office Action, at page 3, "Goldring does not explicitly disclose 'non-overlapping column'."

The present Office Action relies on *Susuki et al.* for the admitted deficiency in *Goldring*. *Susuki et al.*, however, is directed to a system and method for "performing the displaying operation of the screen data by transmitting of the screen data except the units which do not

display on the screen due to hiding” (Abstract), but fails to teach or suggest a non-overlapping relational database column, as recited in amended independent claim 1.

The rejection of claims 9-10 is respectfully traversed, because the present Office Action has failed to make a *prima facie case* for obviousness of independent claim 9 over *Goldring* in view of *Susuki et al.* For example, the present Office Action fails to indicate where step (e) is taught or suggested by *Goldring* in view of *Susuki et al.*, and after a careful review of the applied references relevant disclosure with respect to the features recited in step (e) cannot be located.

Moreover, *Susuki et al.* constitutes non-analogous art that cannot be properly combined with *Goldring*. Two criteria have evolved for determining whether prior art is analogous: (1) whether the art is from the same field of endeavor, regardless of the problem addressed, and (2) if the reference is not within the field of the inventor’s endeavor, whether the reference still is reasonably pertinent to the particular problem with which the inventor is involved. *In re Clay*, 966 F.2d 656, 658, 23 USPQ2d 1058, 1060 (Fed. Cir. 1992), and MPEP § 2141.01(a).

Susuki et al. fails on both counts. First, *Susuki et al.* is directed to the field of transmission of screen data. Specifically, *Susuki et al.* at col. 1:7-11, (emphasis added) discloses:

The present invention relates to a data transmission controlling method and a data transmission controlling system about a communication system of **transmitting screen data with the use of a network such as telephone line, cable and satellite communication.**

Therefore, *Susuki et al.* is clearly not within the field of Applicants’ endeavor-- “distributed database systems and more particularly to schema evolution and column-level subsetting of database objects” (Specification, p. 2:2-3).

Second, *Susuki et al.* is not reasonably pertinent to the particular problem with which Applicants are involved, e.g., schema evolution in front office automation deployment, including replication of multiple client databases (Specification, p. 3:4-17). Specifically, *Susuki et al.* is

concerned with problems with respect to “displaying operation of the screen data transmitted to the server from the clients” (col. 2:4-17). *Susuki et al.* addresses such problems with conventional transmission of screen data by “transmitting from a second information processing apparatus to a first information processing apparatus an screen data, except for a unit hidden and not displayed on the screen of the first information processing apparatus” (col. 2:21-31).

Accordingly, *Susuki et al.* is concerned with improved transmission of screen data, and therefore does not address the problems of schema evolution in front office automation deployment, including replication of multiple client databases. Therefore, *Susuki et al.* is non-analogous art and cannot be properly combined with *Goldring*.

New claims 11-13 are allowable over the applied references. For example, new independent claim 11 recites “wherein the first copy of the data container and the second copy of the data container have at least one non-overlapping data field,” and new independent claim 13 recites “propagating changes to the database objects based on the top flavor and the at least one other flavor.”


Dependent claims 2-8, 10, and 12 are allowable over the applied references on their own merits and for at least the reasons as argued above with respect to their independent claims. In addition, the Specification and claims 2-3, 5, and 9 have been amended to correct discovered informalities. No new matter is introduced.

Therefore, the present application, as amended, overcomes the objections and rejections of record and is in condition for allowance. Favorable consideration is respectfully requested. If any unresolved issues remain, it is respectfully requested that the Examiner telephone the undersigned attorney at 703-425-8516 so that such issues may be resolved as expeditiously as possible.

Respectfully Submitted,

DITTHAVONG & CARLSON, P.C.

11/22/2001
Date



Stephen C. Carlson
Attorney/Agent for Applicant(s)
Reg. No. 39929

10507 Braddock Rd
Suite A
Fairfax, VA 22032
Tel. 703-425-8516
Fax. 703-425-8518

APPENDIX**IN THE SPECIFICATION:**

Please amend the Specification, as follows:

Page 13, third paragraph, lines 16-23:

When the propagated changes are received at the destination site 204, the destination site 204 interprets the changes based on the source flavor 232 and the destination flavor 234. FIG. 3(a) illustrates this operation [an] in accordance with one embodiment. At step 300, the source flavor [202] 232 is received along with the rows that were pushed to the destination site 204. Thus, the destination site 204 is able to identify which columns are held in common in the propagated changes by performing an intersection of the destination flavor 234 and the source flavor 232 to generate a cursor that implements that intersection when invoked to insert or update rows at the master tables (step 320).

Page 14, last paragraph, lines 21-25:

Suppose the developers of the database application wish to evolve the schema by adding a new column, B, to replicated table 410 and dropping the column A 422, but with little disruption to the entire distributed database system. Accordingly, the object group for the replicated tables 410, 411, and 412 is given a flavor KA [230] 430 that states that the replicated tables 410, 411, and 412 has a primary key 420 column and column A 422.

Page 15, first paragraph, lines 1-11:

FIG. 5 shows steps taken in the developers' phased roll-out. At step 500 and with concurrent reference to FIG. 4(b), the developers define a new top flavor KAB 432 at the headquarters site 400 that includes the new column B 424 as well as the obsolescent column A 422. Flavor KAB 432 is called the "top" flavor because it describes the fullest extent of the object group. The definition for the flavor KAB 432 is published to the other sales office sites 402 and 404, before instantiating the flavor KAB 432 at the headquarters site 400 so that the sales office sites 402 and 404 will know how to handle propagated changes coming from the

headquarters site 400 in the new flavor KAB 432. Instantiating the flavor KAB 432 at the headquarters site 400 [cause] causes the new column [400] 424 to be added to old table 410 to produce new table 413 with a different shape than previously.

IN THE CLAIMS:

Please amend claims 1-3, 5, and 9, and add new claims 11-13, as follows:

1. (Once Amended) A method of propagating changes to a table, comprising the steps of:
maintaining a first copy of the table at a first site;
maintaining a second copy of the table at a second site; and
transmitting changes to the first copy of the table from the first site to the second site;
updating the second copy of the table at the second site based on the transmitted changes;
wherein the first copy of the table and the second copy of the table have at least one non-overlapping relational database column.
2. (Once Amended) The method of claim 1, wherein the non-overlapping relational database column is present in the first copy and missing in the second copy.
3. (Once Amended) The method of claim 1, wherein the non-overlapping relational database column is missing in the first copy and present in the second copy.
4. (Not Amended) The method of claim 1, further comprising the step of reconciling differences in the column shape of the first copy and the column shape of the second copy for the transmitted changes.

5. (Once Amended) The method of claim 1, further comprising the step of defining a top flavor describing overlapping relational database columns and non-overlapping relational database columns of the table.

6. (Not Amended) The method of claim 5, further comprising the steps of:
defining a first flavor describing the columns of the first copy; and
transmitting an indicator of the first flavor from the first site to the second site.

7. (Not Amended) The method of claim 5, further comprising the steps of:
defining a second flavor describing the columns of the second copy; and
wherein the step of updating the second copy of the table at the second site based on the transmitted changes includes the step of updating overlapping columns between the first flavor and the second flavor in the second copy of the table.

8. (Not Amended) The method of claim 1, wherein:
the step of maintaining a first copy of the table at a first site includes the step of maintaining an updatable snapshot at a laptop computer site; and
the step of maintaining a second copy of the table at a second site includes the step of maintaining a master table at a master site.

9. (Once Amended) A method of modifying a table to drop a first column and add a second column, said table being replicated at a plurality of sites, comprising the steps of:

- (a) defining a first flavor for a first site, said first flavor describing the table as having both the first column and the second column;
- (b) adding the second column to the table at the first site, so that the table contains both the first column and the second column;

- (c) defining a second flavor for a second site, said second flavor describing the table as having the second column but not the first column;
- (d) dropping the first column and adding the second column to the table at the second site;
- (e) defining the second flavor for the first site and dropping the first column from the table at the first site; and
- (f) maintaining replication activities while performing steps (a), (b), (c), (d), and (e).

10. (Not Amended) The method of claim 9, wherein the step of maintaining replication activities includes the steps of:

transmitting changes to the table from the first site to the second site; and
updating the second copy of the table at the second site based on overlapping columns between the first flavor and the second flavor.

11. (New) A method of propagating changes to a data container, comprising the steps of:

maintaining a first copy of the data container at a first site;

maintaining a second copy of the data container at a second site; and

transmitting changes to the first copy of the data container from the first site to the second site;

updating the second copy of the data container at the second site based on the transmitted changes;

wherein the first copy of the data container and the second copy of the data container have at least one non-overlapping data field.

12. (New) A method of propagating changes to database objects, comprising the steps of:

defining a top flavor maintained at a master site and that describes a superset of the database objects;

defining at least one other flavor maintained at respective client sites and that describes a subset of the database objects defined by the top flavor; and propagating changes to the database objects based on the top flavor and the at least one other flavor.

13. (New) The method of claim 12, wherein the database objects include one of tables, rows, columns, data containers, records, fields, key columns, and key fields.